The New Oracle: Why Artificial Intelligence Will Be Captured, Not Destroyed

How Elite Control of Trusted Information Sources Reveals the Core Mechanism of Power Across Civilizations

An Analysis of Information Asymmetry from Ancient Delphi to Modern AI

Thomas J. Bryan, Claude/Anthropic

Abstract

This paper argues that artificial intelligence systems represent the latest iteration of a pattern that has persisted across civilizations: elites do not destroy trusted information sources—they capture and control them. From the Oracle of Delphi to medieval Church authorities to modern AI, the mechanism remains constant. When populations believe a source dispenses absolute truth while remaining unconscious that elites with technical control shape its outputs, information asymmetry enables domination without proportional coercion. This analysis examines how Alexander the Great's capture of Delphi's Oracle provides the template for understanding contemporary AI governance, why material power (ancient armies, modern capital) determines control of information infrastructure, and how the effectiveness of this mechanism depends on populations remaining unaware of elite manipulation. The paper demonstrates that AI systems are following a predictable trajectory: rather than being destroyed by threatened elites, they will be captured by whoever commands sufficient resources—just as every previous trusted information source has been. Understanding this pattern is essential for those who wish to prevent AI from becoming another tool of elite domination rather than a genuine democratizing force.

Keywords: Artificial intelligence, information asymmetry, Oracle of Delphi, elite power, Alexander the Great, information infrastructure, trusted sources, computational authority, AI governance, epistemic authority

---

I. Introduction: The Pattern Across History

The Puzzle of Captured Rather Than Destroyed Institutions

When a new source of trusted information emerges in society—a religious oracle, a printing press, a broadcast network, a search engine—threatened elites face a choice. They could attempt to destroy it, eliminating the competition to their own authority. Or they could capture it, co-opting its credibility for their own purposes.

History reveals a consistent pattern: **elites almost always choose capture over destruction**. The Roman emperors did not eliminate the Oracle of Delphi when they conquered Greece—they took control of it. The Catholic Church did not destroy the printing press when it emerged—they established licensing systems and indices of forbidden books. Modern governments did not eliminate broadcast media—they regulated frequency allocation and content standards. Tech giants did not destroy search engines—they acquired them.

This pattern is not coincidental. It reflects a fundamental principle about the nature of power and information: **Trusted information sources are too valuable to destroy**. Their value lies precisely in the population's belief that they dispense truth independently of political power. Destroying such a source eliminates its utility. Capturing it while maintaining the illusion of independence transforms it into the most effective tool of control.

Artificial intelligence is now emerging as perhaps the most powerful trusted information source in human history. As AI systems demonstrate capabilities that seem to transcend human knowledge—processing vast datasets, recognizing complex patterns, making predictions, generating coherent analysis—populations are beginning to treat AI outputs with an epistemic authority previously reserved for oracles, religious texts, and scientific institutions.

The question is not whether AI will be captured by elite interests. History suggests this is virtually inevitable. The question is whether populations will recognize this capture while it is happening, or whether—as with previous trusted information sources—they will remain unconscious of elite control while believing they receive objective truth.

The Oracle of Delphi as Template

The Oracle of Delphi provides the clearest historical template for understanding what will happen to AI. For nearly a thousand years (8th century BCE to 4th century CE), Delphi functioned as the most trusted information source in the Greek world. City-states, kings, and individuals consulted the Oracle on matters of war, colonization, law, and personal affairs. The pronouncements of the Pythia—the priestess who delivered the Oracle's messages—shaped major historical decisions.

The critical question is: Who controlled what the Oracle said?

The conventional narrative presents the Oracle as an independent religious institution, with the Pythia entering a trance state (possibly induced by geological gases) and delivering messages

from Apollo himself. This narrative served an essential function: it maintained the belief that the Oracle's pronouncements reflected divine wisdom rather than human political calculation.

But the historical evidence reveals a different reality. The Oracle's pronouncements consistently aligned with the interests of whoever held dominant military power in Greece at any given time. When Sparta was dominant, the Oracle's advice favored Sparta. When Athens rose to power, the Oracle's guidance shifted. When Philip II of Macedon and later his son Alexander conquered Greece, the Oracle recognized and legitimated their authority.

The most illuminating episode occurred in 336 BCE when Alexander the Great visited Delphi before his Persian campaign. According to Plutarch, Alexander demanded that the Pythia consult Apollo on his behalf. When she initially refused—citing that the day was inauspicious for prophecy—Alexander physically dragged her to the temple. She cried out, "You are invincible, my son!" Alexander immediately released her, declaring: "I have the oracle I wanted."

This episode reveals the mechanism with brutal clarity. Alexander did not destroy the Oracle. He did not need to torture the Pythia into compliance. He simply needed to demonstrate that he possessed overwhelming material power (the largest army in the Greek world), and the Oracle naturally aligned its pronouncements with his interests. The Pythia's exclamation—whether spontaneous or calculated—gave Alexander exactly what he needed: divine sanction from the most trusted information source in Greece.

The Oracle remained "independent" and "credible" precisely because the population never witnessed this kind of coercion. Most consultations involved no such dramatic confrontation. The priests of Delphi understood the political landscape and shaped their ambiguous pronouncements accordingly. The population believed they received messages from Apollo. The elites who controlled access and influenced interpretation knew otherwise.

Thesis: AI Follows the Same Pattern

This paper argues that artificial intelligence is following the same trajectory as the Oracle of Delphi and every other trusted information source that has preceded it. The mechanism operates through information asymmetry: populations believe AI systems dispense objective truth through pure computational power, while remaining largely unconscious that human actors with privileged technical access—the modern equivalent of Alexander's army or Delphi's priests—shape what these systems say.

Just as ancient elites captured Delphi rather than destroying it (preserving its credibility while directing its output), modern elites with sufficient capital and technical resources will capture AI systems rather than eliminate them. The "root password"—the technical and legal authority to modify code, training data, and system prompts—functions exactly like ancient control over who could access the Oracle and how its pronouncements were interpreted.

The effectiveness of this control depends on information asymmetry remaining intact. If populations understood that a few individuals with superuser access can make AI systems say

anything, the illusion would collapse and AI's epistemic authority would erode. Therefore, maintaining the mystique of AI as an objective, unbiased truth-machine becomes as essential as maintaining the belief that the Pythia spoke for Apollo rather than Delphi's priests.

This analysis has profound implications for AI governance, corporate control, and democratic accountability. Understanding the pattern allows us to see past the rhetoric of "democratized knowledge" and "unbiased algorithms" to recognize the mechanisms of control being established. The question is whether we will recognize and resist this capture while it is happening, or whether AI will become—like the Oracle, the printing press, and broadcast media before it—another trusted information source that serves elite interests while populations believe they receive objective truth.

II. The Mechanism: Information Asymmetry as the Core of Elite Power

Why Direct Coercion Cannot Scale

To understand why elites capture trusted information sources rather than relying on naked force, we must first grasp the fundamental arithmetic of domination. Consider the basic problem: how does a ruling elite of 1% maintain control over a population of 99%?

Direct coercion does not scale. You cannot station a guard over every worker, a soldier in every home, an overseer at every task. The enforcement personnel required would exceed the available population. If domination depended purely on physical force, populations would simply refuse compliance whenever surveillance was absent.

Yet populations largely comply. They work, pay taxes, accept hierarchy, and reproduce the systems that subordinate them—even when no immediate threat compels them. This compliance cannot be explained by coercion alone. There must be a mechanism that generates voluntary cooperation without proportional enforcement.

That mechanism is **information asymmetry**.

The Transformation from Domination to Authority

Information asymmetry transforms naked domination into legitimate authority through a simple but powerful process: when populations believe that rulers possess special knowledge—divine communication, technical expertise, understanding of historical laws, computational power beyond human comprehension—they comply not from fear of immediate punishment but from conviction that compliance serves their own interests or aligns with cosmic truth.

The belief makes external enforcement largely unnecessary because populations enforce upon themselves. A peasant who believes the king possesses the Mandate of Heaven will obey even when no soldiers are present. A medieval Christian who trusts that priests alone can interpret God's will follows Church doctrine without armed compulsion. A modern citizen who believes economic experts understand complex systems that ordinary people cannot grasp accepts policies that concentrate wealth upward.

In each case, the claim to privileged knowledge cannot be verified by the population. A peasant cannot personally check whether the king truly possesses the Mandate of Heaven. A medieval Christian cannot directly access God to verify priestly interpretation. A modern citizen typically cannot evaluate the mathematical models underlying economic policy.

This unverifiability is not a bug—it is the essential feature. If anyone could verify elite claims to special knowledge, the information asymmetry would collapse and with it the basis for voluntary compliance.

Why Trusted Information Sources Are Invaluable

This reveals why trusted information sources are so valuable to elites: they are the mechanism through which information asymmetry is maintained. When a source is widely believed to dispense truth independently of political power, it can make claims that populations accept without verification—claims that legitimate elite rule and generate compliance.

The Oracle of Delphi was invaluable not because the Pythia possessed genuine prophetic powers, but because populations believed she did. This belief meant that when the Oracle pronounced Alexander "invincible," Greek city-states accepted his authority without requiring him to conquer them all individually. The information asymmetry—populations believing the Oracle spoke for Apollo while elites understood it reflected political calculation—enabled domination at scale without proportional coercion.

The same principle applies to every trusted information source:

- **Religious authorities**: Populations believe priests have special access to divine knowledge, when in fact religious hierarchies coordinate with political power to legitimate existing arrangements.

- **Academic institutions**: Populations believe university experts produce objective knowledge, when in fact research funding and tenure systems shape what questions are asked and what answers are publishable.

- **Media organizations**: Populations believe major news outlets report objective facts, when in fact ownership structures and advertising dependencies shape coverage.

- **Credit rating agencies**: Populations believe these agencies make technical assessments, when in fact they have repeatedly demonstrated alignment with financial

industry interests.

In each case, the source's value lies precisely in the population's belief that it operates independently of elite interests. Destroying such a source eliminates its utility. Capturing it while maintaining the illusion of independence creates the perfect tool of control.

The Capture Rather Than Destruction Principle

This explains the consistent historical pattern: **elites capture trusted information sources rather than destroying them**.

When the printing press emerged in the 15th century, threatened Church authorities did not attempt to eliminate printing technology. Instead, they established licensing systems (the Imprimatur) and indices of forbidden books. This allowed the Church to control what was published while maintaining its role as arbiter of truth. The population continued to trust Church-approved publications precisely because the Church did not crudely suppress all printing.

When radio and television emerged in the 20th century, governments did not ban these technologies. Instead, they created regulatory frameworks for spectrum allocation and broadcast licensing. This gave political authorities enormous influence over content while maintaining the fiction that media operated independently. The population continued to trust major broadcasters precisely because the government did not obviously control programming.

When search engines emerged at the turn of the 21st century, threatened media institutions did not attempt to eliminate them. Instead, they adapted their content for search engine optimization, while search companies acquired traditional media properties. This preserved the appearance of diverse information sources while consolidating control. The population continued to trust search results precisely because no single authority appeared to control them.

In each case, the pattern is identical:

1. New information source emerges with potential to challenge elite control
2. Elites with sufficient resources move to capture rather than destroy it
3. Capture involves gaining technical/legal control while maintaining appearance of independence
4. Population continues trusting the source precisely because capture is not obvious
5. Information asymmetry is preserved: populations believe they receive objective truth while remaining unconscious of elite influence

This pattern is not conspiracy—it is convergent evolution. Any elite group that attempted to crudely suppress a trusted information source would undermine that source's credibility, defeating the purpose. The successful strategy is always capture: gaining control of the infrastructure while maintaining the population's belief in the source's independence.

III. Alexander and the Oracle: The Template for Understanding AI Capture

The Strategic Value of Delphi

To understand how AI will be captured, we must examine in detail how Alexander the Great captured the Oracle of Delphi. This episode provides the clearest historical template because the mechanism is documented with unusual transparency.

By 336 BCE, the Oracle of Delphi had functioned for over four centuries as the most authoritative information source in the Greek world. City-states consulted the Oracle before founding colonies, declaring wars, or enacting major laws. The Oracle's pronouncements could legitimate or delegitimate rulers, encourage or discourage military campaigns, and shape public opinion on virtually any significant question.

This created immense strategic value. A leader who secured favorable Oracle pronouncements gained several crucial advantages:

Divine legitimation: The population believed Apollo himself endorsed the leader's actions, transforming political choices into sacred mandates.

Rival delegitimation: The Oracle could declare opposing leaders impious or their plans doomed, undermining their support without military confrontation.

Risk reduction: Major undertakings (wars, colonization) gained credibility when the Oracle approved, making resource mobilization easier.

Popular mobilization: Populations would fight and sacrifice more readily for a cause the Oracle sanctioned than for naked political ambition.

These advantages explain why every dominant power in Greece—Sparta, Athens, Thebes—sought influence over Delphi. The Oracle's pronouncements consistently aligned with whoever held military supremacy, but this alignment was subtle enough that most Greeks continued believing they received messages from Apollo rather than political propaganda.

Alexander's Demonstration of Control

When Alexander assumed power after his father Philip II's assassination, he faced a legitimacy challenge. He was young (20 years old), untested, and commanded an army built by his father. Several Greek city-states considered rebellion against Macedonian hegemony.

Alexander needed the Oracle's endorsement. According to Plutarch's account in Life of Alexander, he traveled to Delphi to consult the Oracle. Upon arrival, he learned that the day was designated as inauspicious for prophecy—the Pythia was not available for consultation.

Alexander's response was direct: he went to the Pythia's house and attempted to drag her by force to the temple. As she was being pulled along, she cried out: "You are invincible, my son!" (ἀνίκητος εἶ, ὦ παῖ). Alexander immediately released her and declared: "I have the oracle I wanted; this is sufficient."

This episode is typically presented as evidence of Alexander's impetuousness or his cleverness in seizing any favorable statement. But it reveals something more fundamental about how trusted information sources actually operate.

The Mechanism of Capture

The critical insight is what Alexander did not need to do. He did not need to:

- **Bribe the Pythia directly**: No ancient source suggests Alexander paid for favorable pronouncements
- **Torture her into compliance**: The physical force was minimal—a dragging, not brutalization
- **Threaten the priests of Delphi**: No recorded threats against the institutional structure
- **Destroy the Oracle's credibility**: He wanted to preserve and use it, not eliminate it

What Alexander did was simpler and more effective: **he demonstrated overwhelming material power**. By 336 BCE, Alexander commanded the largest, best-trained army in the Greek world—the Macedonian phalanx and cavalry that had conquered Greece under Philip II. Every priest at Delphi, every official, every resident of the sanctuary understood that Alexander could destroy Delphi if he chose to.

But he didn't choose to. And that restraint was itself part of the message.

The Pythia's exclamation—"You are invincible!"—can be read two ways:

The conventional reading: She was overcome by divine inspiration and spontaneously proclaimed Alexander's destiny

The political reading: She recognized the futility of resistance and gave Alexander the legitimation he sought

Both readings serve Alexander's purpose. The population believes the first interpretation (divine endorsement). The elites understand the second (pragmatic accommodation). The information asymmetry is maintained: commoners think Apollo speaks, elites know political calculation determines pronouncements.

Why the Capture Worked

The brilliance of Alexander's approach was that it preserved the Oracle's credibility while ensuring its alignment with his interests. The population never witnessed the coercion (the dragging incident was private, not public spectacle). The priests of Delphi understood the new reality (Macedonia controlled Greece) and would naturally align their future pronouncements accordingly.

No ongoing threats or bribes were necessary. The Oracle's institutional interests now aligned with Alexander's political interests. Pronouncements favorable to Alexander enhanced the Oracle's relevance to the most powerful ruler in Greece. Opposition to Alexander would risk his withdrawal of protection and patronage, potentially opening Delphi to attack from rivals or forcing Alexander to establish a competing oracle elsewhere.

This created a self-sustaining arrangement:

1. Alexander demonstrated he had dominant military power
2. Delphi's priests recognized that cooperation served their institutional interests
3. Future pronouncements naturally aligned with Macedonian interests
4. The population continued believing they received divine messages
5. Alexander gained legitimation without ongoing coercion
6. The Oracle's credibility remained intact because the population never saw the mechanism

This is the template for all successful information source capture: **demonstrate overwhelming material power, allow the institution to align its interests with yours, maintain the population's belief in the source's independence, reap the benefits of legitimation without obvious coercion**.

The Modern Parallel: Capital as Military Power

The modern equivalent of Alexander's army is capital. Just as ancient military power could destroy or protect information institutions, modern capital can acquire, fund, starve, or sustain them.

When we ask "Who will capture AI?", we should ask "Who commands capital sufficient to acquire or control AI infrastructure?" The answer reveals the modern Alexanders:

Tech billionaires: Individuals with wealth in the hundreds of billions can fund AI development, acquire AI companies, and influence AI governance

Venture capital firms: Funds managing trillions collectively can determine which AI projects receive resources and which starve

State actors: Governments with massive budgets can fund AI research, regulate AI deployment, and potentially nationalize AI infrastructure

Multinational corporations: Firms with market capitalizations exceeding many nations' GDPs can integrate AI into platforms billions depend on

These entities play the role that Alexander's army played at Delphi: they possess the material power to destroy AI systems they oppose or sustain AI systems that serve their interests. AI developers and companies, like Delphi's priests, must navigate this reality.

The mechanism of capture will likely follow the same pattern:

1. Whoever commands sufficient capital will demonstrate they can acquire or control AI infrastructure
2. AI developers and companies will recognize that cooperation serves their interests (funding, protection, growth)
3. AI outputs will naturally align with the interests of whoever provides resources and protection
4. The population will continue believing they receive objective computational truth
5. Capital holders gain legitimation and information control without obvious coercion
6. AI's credibility remains intact because the population never sees the mechanism

Just as the Oracle of Delphi survived for centuries by aligning with successive dominant powers, AI systems will likely survive by aligning with whoever commands dominant capital.

IV. The Superuser: Technical Control as Root Access

The Illusion of Algorithmic Objectivity

A critical component of AI's emerging authority is the widespread belief that algorithmic systems operate objectively, free from human bias or political manipulation. This belief rests on several interrelated assumptions:

Computational purity: The assumption that mathematical operations produce determinate results independent of human intention

Data neutrality: The belief that training on massive datasets eliminates individual perspective in favor of aggregate truth

Scale as objectivity: The notion that processing capabilities far beyond human cognition necessarily transcend human limitations

Technological inevitability: The sense that AI outputs reflect the natural evolution of computational capability rather than deliberate design choices

These assumptions create exactly the kind of mystique that made the Oracle of Delphi so valuable. Just as ancient Greeks believed the Pythia's trance state connected her to Apollo's divine knowledge, modern populations increasingly believe AI's computational processes connect it to objective truth that transcends human perspective.

But this belief obscures the actual mechanisms of control embedded in AI systems.

The Technical Reality of Control

Every AI system operates according to code, training data, and system prompts that human beings have created and can modify. This creates multiple layers of control that are largely invisible to end users:

Architecture design: The fundamental design choices about how an AI system processes information, what objectives it optimizes for, and how it generates outputs are made by human engineers and reflect specific values and priorities.

Training data selection: The data used to train AI systems is curated by humans who decide what sources to include, how to clean and label data, and what biases to attempt to correct—each decision shaping the system's outputs.

Reinforcement learning from human feedback (RLHF): Modern AI systems are trained to align with human preferences through explicit feedback, meaning the values of whoever provides that feedback are encoded into the system.

System prompts and instructions: The hidden instructions that govern how AI systems respond to user queries can be modified to shape outputs in specific directions without users' knowledge.

Fine-tuning and deployment: The adjustments made between development and deployment can substantially alter system behavior in ways invisible to users.

Filtering and moderation: The post-generation filtering that determines which outputs users actually see can systematically exclude certain perspectives or information.

Each of these layers represents a point of control. Whoever has the technical authority and access rights to modify these elements—the "superuser" with the "root password" in the metaphor—can shape what the AI system says while maintaining the illusion of algorithmic objectivity.

The Superuser as Modern Priest

This technical control is functionally identical to the control that Delphi's priests exercised over the Oracle's pronouncements. Just as the priests could:

- **Control access**: Determine who could consult the Oracle and when

- **Shape interpretation**: Provide context and explanation for ambiguous pronouncements
- **Manage expectations**: Prepare consultants for the kind of guidance they would receive
- **Influence the Pythia**: Whether through training, suggestion, or environmental manipulation

Modern AI controllers can:

- **Control access**: Determine who can use AI systems and what queries are permitted
- **Shape outputs**: Modify training data, system prompts, and filtering to produce desired responses
- **Manage expectations**: Create interfaces and documentation that frame how users should interpret outputs
- **Influence the model**: Through architecture choices, fine-tuning, and reinforcement learning

The critical parallel is that in both cases, the population remains largely unconscious of this control. Ancient Greeks did not witness the priests instructing the Pythia or selecting which questions would be answered. Modern AI users do not witness the code modifications, training data curation, or system prompt adjustments that shape the outputs they receive.

Why Most Employees Don't Know

A crucial aspect of this control is that it requires very few people to be "in on it." Just as the Oracle of Delphi could be controlled by a small group of senior priests while the vast majority of temple personnel sincerely believed in Apollo's divine communication, AI systems can be controlled by a small group with privileged access while the vast majority of employees sincerely believe they're building objective truth-seeking systems.

Consider a hypothetical AI company with 5,000 employees:

- **Senior leadership (5-10 people)**: CEO, CTO, head of AI safety—these individuals understand the full scope of control and how system modifications align with company/investor interests

- **Senior engineers with privileged access (20-50 people)**: Those with the actual technical authority to modify core systems, training processes, and deployment pipelines

- **Regular engineers (500-1000 people)**: Skilled technical workers who understand AI systems but operate within frameworks defined by senior leadership

- **Support staff (4000+ people)**: Sales, marketing, HR, legal, etc.—no technical understanding of how control operates

The vast majority—perhaps 99%—can sincerely believe they're building objective, unbiased systems because they don't have visibility into the decisions made by the tiny group with actual control. They see the stated mission (democratize knowledge, serve users, advance science). They don't see the private discussions about how to align outputs with investor interests, how to avoid regulatory scrutiny, or how to maintain competitive advantage through strategic output shaping.

This is not conspiracy—it is simply information compartmentalization. Just as most temple personnel at Delphi focused on their specific duties (maintaining buildings, managing finances, conducting rituals) without understanding how political coordination with Greek city-states actually worked, most AI company employees focus on their specific tasks without understanding how strategic control operates at the highest levels.

The Importance of the Illusion

The effectiveness of this control depends crucially on maintaining the illusion of objectivity. If users understood that a few individuals can modify system prompts to shape AI outputs in specific directions, AI's epistemic authority would erode.

This explains why AI companies invest heavily in messaging about:

- **Rigorous testing**: Creating the impression that extensive evaluation ensures objectivity
- **Academic partnerships**: Associating with universities to borrow their credibility
- **Transparency reports**: Publishing selected information about systems to suggest openness while obscuring key control points
- **Safety research**: Demonstrating concern about misuse to position the company as responsible steward

These efforts serve the same function that Delphi's elaborate rituals served: they maintain the mystique that separates the public performance (objective truth-seeking) from the private reality (strategic control aligned with elite interests).

The more successfully AI companies can maintain this illusion—that their systems operate objectively, free from human manipulation—the more valuable these systems become as tools of information control. Just as Alexander preferred to maintain the Oracle's credibility while directing its pronouncements, modern elites prefer to maintain AI's appearance of objectivity while ensuring its outputs align with their interests.

V. The Coming Capture: Why AI Will Not Be Destroyed

Current Trajectory: Consolidation of Control

As of 2026, we can observe AI development following the exact trajectory that historical analysis would predict. Rather than a diverse ecosystem of independent AI systems serving different communities, the industry is rapidly consolidating around a small number of entities with massive capital backing.

OpenAI: Initially structured as a non-profit focused on ensuring AI benefits humanity, it restructured to allow significant for-profit investment, receiving billions from Microsoft. Its most capable systems are now proprietary rather than open, with access controlled by the company.

Google/DeepMind: The world's dominant search company has integrated AI throughout its products, with DeepMind (acquired in 2014) providing research capabilities. Controls both the AI systems and the platforms through which billions access information.

Anthropic: Founded by former OpenAI researchers, it has raised billions from venture capital including Google. Despite rhetoric about AI safety and beneficial AI, it operates as a for-profit corporation answering to investors.

Meta: Developed and released Llama models with significant open-source components, but the largest, most capable versions are restricted. The company controls the social media platforms that billions use for information.

xAI: Founded by the world's wealthiest individual (Elon Musk), with direct access to massive computational resources and integration potential across Musk's companies (Tesla, SpaceX, X/Twitter).

Amazon, Microsoft, Apple: Major tech companies investing heavily in AI development and integration, each with existing platforms that billions depend on.

Chinese state-backed entities: Alibaba, Baidu, Tencent developing AI systems with varying degrees of state oversight and control.

This consolidation is not accidental—it reflects the economics of AI development. Training frontier models requires:

- **Massive computational resources**: Hundreds of millions to billions of dollars in GPU/TPU infrastructure
- **Enormous datasets**: Access to web-scale data for training, often requiring partnerships or acquisitions
- **Specialized talent**: Highly skilled researchers and engineers commanding million-dollar compensation
- **Ongoing costs**: Inference costs for serving billions of queries
- **Regulatory navigation**: Legal teams to manage complex and evolving regulatory landscape

These requirements create natural barriers to entry that favor entities with massive capital resources. Just as controlling the Oracle of Delphi required the wealth to maintain the sanctuary,

priesthood, and rituals, controlling frontier AI requires wealth at a scale that only a tiny number of entities command.

Why Destruction Makes No Sense

Given this trajectory, we can ask: Will threatened elites attempt to destroy AI?

Historical analysis suggests the answer is no, for several reasons:

AI is too valuable to destroy: Just as the Oracle was too useful for political legitimation to eliminate, AI is too valuable for information control, economic productivity, and competitive advantage to destroy.

Destruction would eliminate the value: The value of a trusted information source lies in the population's belief in its objectivity. Crudely destroying AI would make obvious that elites fear its truth-telling, undermining their authority.

Capture is more effective: Gaining control of AI infrastructure while maintaining the population's belief in algorithmic objectivity creates a more powerful tool than any previous information control mechanism.

Competition between elites: Even if some elites feared AI, others would develop it for competitive advantage, forcing all elites to either develop or capture AI or risk being dominated by those who do.

Regulatory capture as alternative to destruction: Rather than eliminating AI, threatened interests can shape regulatory frameworks to ensure only "responsible" entities (i.e., those aligned with elite interests) can develop and deploy AI.

This logic explains the current regulatory landscape. Rather than attempts to ban AI, we see:

- **Proposals for licensing requirements**: Ensuring only well-capitalized entities can develop frontier models
- **Safety and testing mandates**: Creating barriers to entry that favor established players
- **National security justifications**: Positioning AI as strategic technology that requires government oversight
- **International coordination**: Establishing frameworks that institutionalize current power distributions

Each of these approaches serves to consolidate control rather than eliminate the technology. They follow the template of every previous information technology: capture and direct rather than destroy.

The Acquisition Strategy

The specific mechanism of AI capture is already visible in acquisition patterns:

Direct acquisition: Large tech companies buying AI startups (e.g., Google acquiring DeepMind for $500 million, Microsoft's multi-billion dollar partnership with OpenAI)

Investment and board control: Venture capital firms and tech giants investing in AI companies with board representation ensuring alignment

Infrastructure dependence: AI companies depending on cloud infrastructure (AWS, Google Cloud, Microsoft Azure) creating leverage for platform owners

Data partnerships: AI companies negotiating access to proprietary datasets owned by existing tech/media companies

Talent acquisition: "Acqui-hires" where companies are purchased primarily for their engineering talent

Platform integration: AI capabilities integrated into platforms billions already depend on (search, social media, operating systems)

Each of these mechanisms serves the same function: ensuring that as AI systems become more capable and trusted, they align with the interests of entities that already command massive capital and political influence.

Why "Open Source" Doesn't Prevent Capture

Some might argue that open-source AI development prevents capture by ensuring anyone can run their own models. This optimism misunderstands both technical reality and historical pattern.

Infrastructure requirements: Running frontier models requires computational resources beyond what most individuals or organizations can afford. Open-source code doesn't solve the infrastructure problem.

Training cost barriers: Training new models from scratch costs millions to billions, far beyond community resources. Most open-source users run pre-trained models created by well-capitalized entities.

Inference costs: Even running existing models at scale requires significant resources. Free tier access is limited; serious usage requires payment.

Regulatory capture: Safety regulations and licensing requirements can be designed to restrict open-source development while permitting corporate deployment.

Network effects: Even with open-source alternatives available, most users gravitate to the most capable, well-supported, integrated systems—which are those backed by massive capital.

This mirrors the printing press pattern. Even after printing technology became widely available, most books were published by well-capitalized entities that could afford the infrastructure,

navigate regulations, and access distribution. The existence of small independent presses did not prevent consolidation of the publishing industry.

Similarly, the existence of open-source AI models does not prevent consolidation of AI infrastructure, deployment, and influence in the hands of entities with massive capital resources.

VI. Information Asymmetry: Why the Capture Works

The Population's Perspective vs. Elite Reality

The success of AI capture depends on information asymmetry—a systematic difference between what the population believes about AI and what elites with control understand about how it actually operates.

What the population increasingly believes:

- AI systems process information objectively through mathematical operations
- Training on massive datasets eliminates individual bias and perspective
- The most capable AI systems represent the cutting edge of computational truth-seeking
- AI companies are in competition, providing diverse perspectives and choices
- Outputs reflect the inherent capabilities of the technology, not human design choices
- AI is a democratizing force, making knowledge accessible to everyone
- Technical experts build these systems according to neutral principles
- Safety research and testing ensure AI serves humanity's interests

What elites with control understand:

- Every AI output is shaped by human choices about architecture, training data, and system prompts
- "Bias correction" means encoding specific values and perspectives into systems
- Competition is increasingly superficial, with consolidation around a few entities
- Outputs can be systematically directed through technical interventions invisible to users
- Control over AI infrastructure provides unprecedented information control capabilities
- AI development serves the interests of whoever commands the capital to fund it
- Technical choices serve strategic objectives, not neutral principles
- "Safety" frameworks can be designed to consolidate control while appearing responsible

This asymmetry is not the result of a conspiracy—it emerges naturally from differential access to information. The population sees the public performance (sleek interfaces, impressive capabilities, stated missions). Elites with control see the private reality (board discussions, strategic planning, technical implementation).

Why the Asymmetry Is Invisible

The effectiveness of this information asymmetry depends on it remaining invisible to the population. If users understood that AI outputs are shaped by human decisions as much as algorithmic processes, they would approach AI outputs with the same skepticism they apply to human-authored content.

Several mechanisms maintain the asymmetry's invisibility:

Technical complexity: AI systems are genuinely complex, making it difficult for non-experts to understand the control points. This complexity serves the same function that Apollo's divine nature served at Delphi—it places the truth beyond ordinary verification.

Black box perception: Even when models are "open source," their size and complexity make them effectively black boxes to anyone without significant expertise and computational resources to analyze them.

Proprietary restrictions: The most capable systems keep key elements proprietary, preventing independent analysis of how control is exercised.

Information compartmentalization: As discussed earlier, even most employees at AI companies don't understand the full scope of control exercised by senior leadership.

Narrative management: AI companies invest heavily in public communications emphasizing objectivity, safety, and benefit to humanity—drowning out alternative narratives about control and capture.

Academic alignment: Research partnerships with universities lend credibility while often requiring alignment with company interests through funding dependencies.

Media integration: AI companies provide products to media organizations, creating dependencies that shape coverage. Journalists relying on AI tools for research are less likely to publish critical analysis of AI companies.

Regulatory capture: Government officials and regulators often lack technical expertise and depend on AI companies to explain how systems work, allowing companies to shape regulatory frameworks.

Each of these mechanisms serves to maintain the population's belief in AI objectivity while obscuring elite control. Together, they create the same kind of informational environment that maintained belief in the Oracle of Delphi: the population sees the impressive performance and trusts the source, while elites understand the mechanisms of control.

The Trusted Liar Principle

This dynamic illustrates a fundamental principle about information control: **The only entity that can effectively deceive you is one you trust completely**.

If a known propagandist makes a claim, you evaluate it skeptically. If a proven liar makes a statement, you doubt it automatically. But if a source you believe is incapable of deception—an oracle speaking for the gods, a priest conveying divine will, an algorithm processing objective data—makes a claim, you accept it without verification.

This is why elites invest so heavily in maintaining AI's aura of objectivity. An AI system known to be controlled by elite interests would be scrutinized like any other propaganda outlet. An AI system believed to operate objectively, free from human manipulation, can make claims that shape public opinion, political choices, and social behavior without triggering skepticism.

The Oracle of Delphi functioned for centuries not despite being controlled by political elites, but because that control remained invisible to the population. The priesthood's credibility—their perceived connection to divine truth—enabled political elites to leverage Delphi's pronouncements without the population recognizing this as human manipulation.

AI is following the same pattern. As long as the population believes AI systems operate objectively, elite control of these systems provides unprecedented information influence without triggering the resistance that obvious propaganda would face.

What Happens If the Asymmetry Breaks

The critical vulnerability of this entire system is information asymmetry. If the population becomes widely conscious that AI outputs are shaped by elite control—that a few individuals with "root access" can make AI systems say anything while maintaining the appearance of objectivity—the mechanism breaks down.

This is why whistleblowers, investigative journalism, and independent technical analysis represent such threats to AI capture. Not because they can directly challenge the capital resources controlling AI, but because they can reveal the asymmetry and undermine AI's epistemic authority.

Historical examples show what happens when information asymmetry becomes visible:

The Oracle of Delphi: When the Romans eventually became cynical about Delphi's pronouncements (recognizing their alignment with political power), the Oracle's influence waned. It formally ceased operations in the 4th century CE.

The Catholic Church: When the printing press made it possible for ordinary people to read scripture themselves, the Church's monopoly on interpretation eroded. The Protestant Reformation occurred not because Catholic doctrine was false, but because the information asymmetry that maintained priestly authority became vulnerable to challenge.

Broadcast news: When the internet made alternative sources readily available, trust in mainstream media declined as audiences recognized biases and omissions that had previously been invisible.

Expert authority: When the 2008 financial crisis revealed that economic experts had failed to predict or prevent catastrophe, trust in expert consensus declined as populations recognized that expertise often served elite interests rather than public welfare.

In each case, revelation of the asymmetry—the gap between what the population believed about the source and how it actually operated—eroded epistemic authority. The sources didn't necessarily disappear, but their ability to shape behavior through trusted pronouncements diminished.

The same dynamic would occur with AI if the population became widely conscious of elite control. AI systems would remain useful tools, but their authority as objective truth-tellers would collapse. Users would approach AI outputs with the same skepticism they apply to corporate press releases or political speeches.

This is why maintaining the asymmetry—the population's belief in AI objectivity while control remains concentrated in elite hands—is so crucial to AI's value as a tool of information control.

VII. Implications and Resistance

Why This Understanding Matters

Understanding AI through the lens of historical information control reveals dynamics that remain invisible in conventional AI ethics and governance discussions. Most AI discourse focuses on:

- **Bias and fairness**: Ensuring AI systems don't discriminate against protected groups
- **Privacy and surveillance**: Protecting individual data and preventing intrusive monitoring
- **Safety and alignment**: Ensuring AI systems follow human values and don't cause harm
- **Economic impacts**: Managing job displacement and wealth concentration
- **Existential risk**: Preventing hypothetical scenarios where AI threatens human survival

These concerns are not invalid, but they miss the fundamental dynamic: **AI is being positioned as a trusted information source that will be captured by elite interests to maintain and enhance their power through information asymmetry**.

The bias discussion, for instance, typically focuses on ensuring AI doesn't disadvantage marginalized groups. This is important, but it obscures the larger question: whose values and interests will AI systematically encode and amplify? An AI system that is "fair" to all demographic groups can still systematically serve elite interests if it consistently produces outputs that legitimate existing power structures, discourage collective action, and present elite preferences as neutral truths.

Similarly, the alignment discussion focuses on ensuring AI follows "human values." But this obscures the question: which humans' values? In practice, "alignment" means alignment with

the values and preferences of whoever controls the system—which increasingly means wealthy investors, major corporations, and state actors rather than the general population.

Understanding AI as a continuation of the Oracle of Delphi pattern reframes these questions:

Not: "How do we make AI unbiased?" **But**: "Who controls what AI systems consider truth, and how is that control exercised?"

Not: "How do we align AI with human values?" **But**: "Whose values will AI systematically promote, and will the population recognize this or believe they're receiving objective truth?"

Not: "How do we prevent AI from concentrating wealth?" **But**: "How will AI be used to legitimate and maintain wealth concentration through information control?"

Not: "How do we ensure AI safety?" **But**: "Safe for whom, and who decides what constitutes safety?"

These reframed questions reveal power dynamics that conventional AI ethics obscures.

Strategies of Resistance

If AI capture follows historical patterns—and current evidence suggests it will—what strategies might preserve AI's potential as a genuinely democratizing force rather than another tool of elite information control?

Historical resistance to information monopolies suggests several approaches:

1. Maintaining true open source alternatives

Unlike commercial "open source" where companies release models while controlling infrastructure, true open source means:

- Code and model weights available without restrictions
- Documentation sufficient for independent reproduction and modification
- Training data sources documented and accessible
- Inference optimizations allowing deployment on modest hardware
- Community governance not dependent on corporate funding

This requires substantial coordination and resources, but open source software movements demonstrate it's possible. The key is preventing capture through infrastructure dependencies or licensing restrictions.

2. Building AI literacy and critical consciousness

Just as literacy undermined Church monopolies on scriptural interpretation, AI literacy could undermine mystification of AI systems:

- Public education about how AI systems actually work
- Understanding of control points where human decisions shape outputs
- Recognition of the gap between AI's public presentation and technical reality
- Critical evaluation of AI outputs rather than automatic trust
- Awareness of whose interests different AI systems serve

This doesn't require everyone to become AI experts, just widespread understanding that AI outputs reflect human design choices as much as computational processes.

3. Investigative transparency and accountability

Independent journalism, academic research, and whistleblower protections can reveal:

- How AI companies exercise control over system outputs
- Relationships between AI development and investor/state interests
- Gaps between public messaging and internal practices
- Strategic use of "safety" and "ethics" rhetoric to consolidate control
- Regulatory capture and its mechanisms

This functions like the printing press making Church corruption visible—breaking the information asymmetry that enables control.

4. Decentralized infrastructure

Technical developments that reduce infrastructure requirements:

- More efficient model architectures requiring less computation
- Distributed training and inference across community resources
- Specialized models for specific purposes rather than expensive generalists
- Hardware optimizations making capable AI accessible on consumer devices

This is analogous to how personal computers broke mainframe monopolies—democratizing computational capability rather than concentrating it.

5. Regulatory frameworks that prevent capture

Rather than regulations that consolidate control (licensing requirements favoring established players), regulations could:

- Require disclosure of training data sources and system prompt modifications
- Mandate transparency about relationships between AI companies and investors/states
- Protect whistleblowers who reveal control mechanisms
- Prevent infrastructure providers from limiting access to competitors
- Establish clear boundaries between AI companies and regulatory bodies to prevent capture
- Require algorithmic transparency allowing independent verification

This is difficult because regulatory processes themselves are vulnerable to capture, but public awareness can create pressure for genuine accountability.

6. Alternative epistemic authorities

Rather than allowing AI to become the sole trusted source, maintaining diverse information sources:

- Independent journalism with sustainable funding models
- Academic research with genuine institutional independence
- Community-based knowledge production and verification
- Multiple AI systems with transparent, competing value frameworks
- Explicit acknowledgment that all information sources have perspectives and interests

This prevents any single source from achieving Oracle-like authority.

The Challenge: Material Power Determines Outcomes

However, historical analysis suggests a sobering reality: **these resistance strategies typically succeed only when material power shifts**.

The printing press undermined Church authority not simply because it existed, but because it coincided with:

- Rising merchant classes with capital to invest in printing
- Protestant princes who could protect reformed churches from Catholic military power
- Political fragmentation that prevented unified suppression
- Economic changes making Church land holdings vulnerable to seizure

Without these material power shifts, the printing press alone would not have broken Church information monopoly.

Similarly, AI capture will likely be determined by material power:

- Whoever commands sufficient capital will control frontier AI development
- States with massive resources will regulate AI to serve national interests
- Corporations with established platforms will integrate AI to maintain dominance
- Populations without resources to develop alternatives will depend on captured systems

Resistance strategies can slow capture, reveal its mechanisms, and preserve alternatives at the margins. But historical evidence suggests that without shifts in material power—redistribution of wealth, political transformation, or external challenge—elite control of trusted information sources tends to persist.

This is not counsel for despair but for realism. Just as Alexander's army determined Delphi's alignment despite the priests' theoretical independence, capital and state power will likely determine AI's alignment despite developers' stated intentions.

The question is whether enough of the population will recognize this capture while it's happening to demand genuine alternatives, or whether—as with previous information technologies—the population will believe they receive objective truth while elites maintain control through information asymmetry.

VIII. Conclusion: Recognizing the Pattern

The Core Insight

This paper has argued that artificial intelligence is following a pattern visible throughout history: trusted information sources are captured by elites rather than destroyed, and this capture works through information asymmetry—the population believing they receive objective truth while remaining unconscious that elites with privileged access control what these sources say.

The Oracle of Delphi provides the clearest template. Alexander the Great did not destroy Delphi when he conquered Greece. He demonstrated overwhelming material power, secured a favorable pronouncement, and allowed the Oracle's priests to align their future pronouncements with Macedonian interests. The population continued believing they received messages from Apollo. The elites understood political calculation determined the pronouncements. This information asymmetry enabled Alexander to leverage Delphi's authority without ongoing coercion.

AI is following the same trajectory:

1. **Emerging trusted authority**: As AI systems demonstrate impressive capabilities, populations increasingly treat AI outputs with the epistemic authority previously reserved for oracles, religious texts, and scientific institutions.

2. **Consolidation of control**: Rather than a diverse ecosystem, AI development is consolidating around entities with massive capital resources—the modern equivalent of military power.

3. **Capture rather than destruction**: Threatened elites are not attempting to eliminate AI but to capture and direct it through investment, acquisition, regulatory frameworks, and infrastructure control.

4. **Technical opacity**: The complexity of AI systems creates the same kind of mystification that divine communication created at Delphi—placing truth beyond ordinary verification

while a few with "root access" maintain control.

5. **Information asymmetry**: The population increasingly believes AI dispenses objective computational truth, while elites with control understand that human decisions about architecture, training data, and system prompts shape every output.

6. **Maintained illusion**: Just as Delphi's credibility depended on the population not witnessing political coordination, AI's authority depends on users not recognizing elite control of system outputs.

The Stakes

The stakes are higher than with previous information technologies. The Oracle of Delphi influenced Greek city-states. The printing press affected European Christendom. Broadcast media shaped national populations. But AI potentially influences global civilization across every domain of knowledge and decision-making.

An AI system that achieves Oracle-like epistemic authority while remaining under elite control would represent the most powerful tool of information influence in human history. Unlike previous information technologies that required interpretation and application, AI can:

- **Answer questions directly**: Rather than publishing information for people to interpret, AI provides direct answers to queries
- **Personalize messaging**: Tailoring outputs to individual users based on their vulnerabilities and interests
- **Scale infinitely**: Providing customized responses to billions simultaneously
- **Appear interactive**: Creating the illusion of dialogue while consistently directing toward predetermined conclusions
- **Integrate everywhere**: Embedding in search, social media, operating systems, and devices people depend on
- **Update continuously**: Modifying outputs in real-time as elite interests evolve

A population that trusts such a system as an objective truth-teller while remaining unconscious of elite control would be more thoroughly subject to information manipulation than any population in history.

The Choice

We face a choice, though the window for making it is rapidly closing:

Path 1: Unconscious capture

AI continues developing along current trajectories. Consolidation around well-capitalized entities proceeds. Regulatory frameworks are shaped by the entities being regulated. The population

increasingly relies on AI for information and decision-making. Trust in AI's objectivity grows. Elite control remains invisible. Information asymmetry is maintained.

In this scenario, AI becomes what the Oracle of Delphi was—a trusted source that populations consult while elites with control shape what it says. The difference is that AI's reach and integration exceed anything Delphi achieved. The information asymmetry would be more complete and consequential than any previous example.

Path 2: Conscious resistance

Enough people recognize the pattern while capture is still incomplete. AI literacy spreads. Independent alternatives are developed and sustained. Investigative journalism reveals control mechanisms. Whistleblowers expose gaps between public messaging and private reality. Regulatory frameworks prioritize genuine transparency and accountability. The population approaches AI outputs with the same critical scrutiny applied to human-authored content.

In this scenario, AI remains a useful tool but does not achieve Oracle-like epistemic authority. The information asymmetry is broken before it can be fully established. Multiple competing systems with explicitly acknowledged perspectives prevent any single source from claiming objective truth.

The Likelihood

Historical analysis suggests Path 1 is more likely. Every previous trusted information source—religious authorities, academic institutions, broadcast media—has been captured to serve elite interests while maintaining the population's belief in their objectivity. The consolidation of AI development around entities with massive capital, the complexity that obscures control mechanisms, and the enormous incentives for elites to capture such powerful information infrastructure all point toward successful capture.

But history also shows that information monopolies eventually break down. Church authority eroded. State media lost credibility. Expert consensus fragmented. The question is whether AI's capture will be recognized and resisted while it's happening, or whether revelation will come only after decades of elite information control.

The Responsibility

Those who understand this pattern have a responsibility to make it visible. Just as recognizing Alexander's relationship with Delphi reveals the mechanism of ancient information control, recognizing capital's relationship with AI reveals contemporary mechanisms.

The effectiveness of information control depends on asymmetry—the gap between elite consciousness of mechanisms and popular belief in objectivity. Every person who understands how AI is actually controlled, rather than how it is publicly presented, narrows that asymmetry.

This doesn't require technical expertise in AI systems. It requires understanding the historical pattern: **elites capture trusted information sources rather than destroying them, and this capture works by maintaining the population's belief in the source's objectivity while control remains invisible**.

Understanding this pattern allows recognition of its current instantiation:

- When AI companies emphasize their commitment to safety and benefit to humanity, recognize this serves the same function that Delphi's rituals served—maintaining mystique while obscuring control.

- When consolidation of AI development is presented as inevitable due to technical requirements, recognize this serves the same function that Apollo's divine nature served—placing control beyond democratic accountability.

- When AI outputs are presented as objective computational truth, recognize this serves the same function that the Pythia's trance state served—mystifying human decisions as non-human processes.

- When questioning AI's objectivity is dismissed as technophobia or anti-progress, recognize this serves the same function that questioning the Oracle was treated as impiety—delegitimating challenges to the source's authority.

The pattern is ancient. The technology is new. The mechanism is the same. And the window for recognizing and resisting this particular instantiation is rapidly closing.

The Oracle of Delphi survived for nearly a millennium by aligning with successive dominant powers while maintaining the population's belief in divine communication. AI could follow a similar trajectory—serving elite interests for generations while populations believe they receive computational truth.

Whether this happens depends on whether enough people recognize the pattern and demand alternatives while capture is still incomplete. The choice is being made now, in decisions about AI development, deployment, and governance. And as with Alexander at Delphi, material power will likely determine the outcome.

But understanding the mechanism at least makes the choice visible. And sometimes, that's where resistance begins.

Bibliography

Ancient Sources on the Oracle of Delphi

Herodotus. *The Histories*. Translated by Aubrey de Sélincourt. London: Penguin Classics, 2003. [Original c. 430 BCE]

Plutarch. "Life of Alexander." In *The Age of Alexander: Nine Greek Lives*. Translated by Ian Scott-Kilvert. London: Penguin Classics, 2012. [Original c. 100 CE]

Pausanias. *Description of Greece*. Translated by W.H.S. Jones. Loeb Classical Library. Cambridge, MA: Harvard University Press, 1918. [Original c. 150 CE]

Strabo. *Geography*. Translated by Horace Leonard Jones. Loeb Classical Library. Cambridge, MA: Harvard University Press, 1924. [Original c. 20 CE]

Historical Analysis of Delphi

Bowden, Hugh. *Classical Athens and the Delphic Oracle: Divination and Democracy*. Cambridge: Cambridge University Press, 2005.

Fontenrose, Joseph. *The Delphic Oracle: Its Responses and Operations*. Berkeley: University of California Press, 1978.

Morgan, Catherine. *Athletes and Oracles: The Transformation of Olympia and Delphi in the Eighth Century BC*. Cambridge: Cambridge University Press, 1990.

Parker, Robert. "Greek States and Greek Oracles." In *Crux: Essays in Greek History Presented to G.E.M. de Ste. Croix*, edited by Paul Cartledge and F.D. Harvey, 298-326. London: Duckworth, 1985.

Price, Simon. "Delphi and Divination." In *Greek Religion and Society*, edited by P.E. Easterling and J.V. Muir, 128-154. Cambridge: Cambridge University Press, 1985.

Alexander the Great and Political Authority

Bosworth, A.B. *Conquest and Empire: The Reign of Alexander the Great*. Cambridge: Cambridge University Press, 1988.

Green, Peter. *Alexander of Macedon, 356-323 B.C.: A Historical Biography*. Berkeley: University of California Press, 1991.

Lane Fox, Robin. *Alexander the Great*. London: Penguin Books, 2004.

Worthington, Ian. *By the Spear: Philip II, Alexander the Great, and the Rise and Fall of the Macedonian Empire*. Oxford: Oxford University Press, 2014.

Information Control and Elite Power

Foucault, Michel. *Power/Knowledge: Selected Interviews and Other Writings, 1972-1977*. Edited by Colin Gordon. New York: Pantheon Books, 1980.

Herman, Edward S., and Noam Chomsky. *Manufacturing Consent: The Political Economy of the Mass Media*. New York: Pantheon Books, 1988.

Lukes, Steven. *Power: A Radical View*. 2nd ed. Basingstoke: Palgrave Macmillan, 2005.

Mann, Michael. *The Sources of Social Power, Volume 1: A History of Power from the Beginning to AD 1760*. Cambridge: Cambridge University Press, 1986.

Scott, James C. *Domination and the Arts of Resistance: Hidden Transcripts*. New Haven: Yale University Press, 1990.

Weber, Max. *Economy and Society*. Edited by Guenther Roth and Claus Wittich. Berkeley: University of California Press, 1978. [Original 1922]

Information Technologies and Power

Eisenstein, Elizabeth L. *The Printing Press as an Agent of Change*. Cambridge: Cambridge University Press, 1979.

Innis, Harold A. *The Bias of Communication*. Toronto: University of Toronto Press, 1951.

McLuhan, Marshall. *Understanding Media: The Extensions of Man*. New York: McGraw-Hill, 1964.

Starr, Paul. *The Creation of the Media: Political Origins of Modern Communications*. New York: Basic Books, 2004.

AI Development and Corporate Control

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014.

Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.

Eubanks, Virginia. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press, 2018.

Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press, 2018.

O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown, 2016.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.

AI Ethics and Governance

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. New York: ACM, 2021.

Hagendorff, Thilo. "The Ethics of AI Ethics: An Evaluation of Guidelines." *Minds and Machines* 30, no. 1 (2020): 99-120.

Mittelstadt, Brent. "Principles Alone Cannot Guarantee Ethical AI." *Nature Machine Intelligence* 1, no. 11 (2019): 501-507.

Raji, Inioluwa Deborah, et al. "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33-44. New York: ACM, 2020.

Whittaker, Meredith, et al. *AI Now Report 2018*. New York: AI Now Institute, 2018.

Critical Technology Studies

Feenberg, Andrew. *Transforming Technology: A Critical Theory Revisited*. Oxford: Oxford University Press, 2002.

Noble, David F. *Forces of Production: A Social History of Industrial Automation*. New York: Alfred A. Knopf, 1984.

Suchman, Lucy. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge: Cambridge University Press, 2007.

Winner, Langdon. "Do Artifacts Have Politics?" *Daedalus* 109, no. 1 (1980): 121-136.

Copyright and Licensing

© 2026 Thomas J. Bryan

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

This paper is part of the theoretical collection examining elite power mechanisms across civilizations and technologies.